

System Architecture and Hardware Design of the CDF XFT Online Track Processor

S. Holm¹, J. R. Dittman¹, J. Freeman¹, R. Klein¹, J. D. Lewis¹, T. M. Shaw¹, T. Wesson¹, C. Ciobanu², J. Gerstenslager², J. Hoftiezer², R. Hughes², M. Johnson², P. Koehn², C. Neu², C. Sanchez², B.L. Winer², K. Bloom³, D. Gerdes³, J.W. Dawson⁴, W.N. Haberichter⁴

¹Fermi National Accelerator Laboratory, ²The Ohio State University,

³University of Michigan, ⁴Argonne National Laboratory

Abstract

A trigger track processor is being designed for CDF Run 2. This processor identifies high momentum ($P_T > 1.5$ GeV/c) charged tracks in the new central outer tracking chamber for the CDF II detector. The design of the track processor, called the eXtremely Fast Tracker (XFT), is highly parallel and handle an input rate of 183 Gbits/sec and output rate of 44 Gbits/sec. The XFT is pipelined and reports the results for a new event every 132ns. The XFT uses three stages, hit classification, segment finding, and segment linking. The pattern recognition algorithms for the three stages are implemented in Programmable Logic Devices (PLDs) which allow for in-situ modification of the algorithm at any time. The PLDs reside on three different types of modules. Prototypes of each of these modules have been designed and built, and are working. An overview of the hardware design and the system architecture are presented.

I. INTRODUCTION

The eXtremely Fast Tracker (XFT), is a trigger track processor that identifies charged tracks in CDF's new Central Open Tracker (COT). The tracks are found in time for the Level 1 trigger decision, and are used for online identification of electrons and muons. This trigger is at the heart of much of the physics that CDF hopes to do in Run II. The XFT is needed to identify high momentum leptons for top, electro-weak, and exotic physics. It is also needed for identification of low-momentum charged tracks for B physics analyses. Without this device the physics potential of CDF II would be severely compromised.

The architecture of the XFT system consists of three main modules along with the medium of transport of data between the three modules. The task of the three main modules is as follows:

- **XFT TDC Mezzanine**[1] module performs hit classification.(TDC – Time to digital converter module)
- **Finder**[2] module performs segment finding within a superlayer of the COT .
- **Linker**[3] module performs segment linking of segments from four superlayers in the COT.

The main task of the three modules are implemented in PLDs which allow for modification of the algorithm at any time.

The three main modules each have a transition module associated with it that performs the task of interfacing to a cable for transmission of data from one module to another. All modules are housed in VME crates which reside in Relay Racks either in the CDF Collision Hall or the 1st Floor Counting Room. The modules along with the medium of transport between these modules will be detailed in this paper.

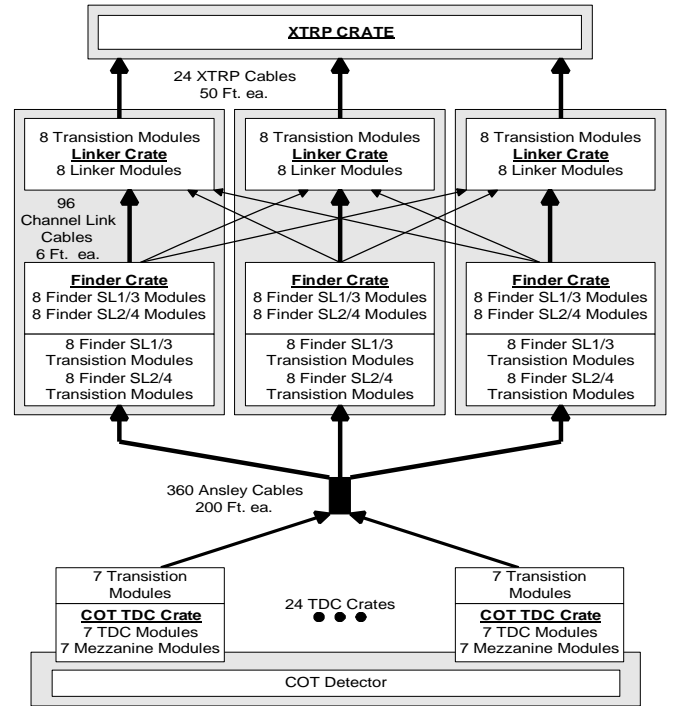


Figure 1: XFT System Diagram.

II. XFT SYSTEM ARCHITECTURE

A block diagram of the XFT System is shown in Figure 1. The XFT System begins with the XFT TDC Mezzanine module residing on the COT TDC's in the CDF collision hall. This module will classify the hits on the COT wires as prompt or delayed and send that information to a transition module at the back of the COT TDC crate. The transition module drives the data at 45.5 MHz, with Low Voltage Differential Signal (LVDS) technology, onto Level 1 Trigger cables that carry the COT wire data 220 ft. to the Finder module crates. Finder Transition modules receive the data and send it across a customized VME backplane to the Finder modules. Finder modules find track segments and report them to Linker modules in another crate. The Finder data is

¹Work supported by the U.S. Department of Energy under contract No. DE-AC02-76CH03000.

transported to the Linker with the use of LVDS Channel Link technology running at ~210 MHz. Linker modules accept segment data from four separate Finder modules. A Linker module uses the Finder data to link together four segments(each from a separate COT superlayer) to form a track. Information from a found track is sent to a Linker Transition module which drives the track information to the XTRP system via LVDS technology running at 15 MHz.

III. XFT TDC MEZZANINE MODULE

The XFT TDC Mezzanine module is a small module that plugs into a TDC module. Throughout the rest of this paper the module will be referred to as the Mezzanine. The Mezzanine module generates trigger primitives that are used to identify hits on wires within COT cells. Each TDC module contains 96 channels and receives discriminated COT signals from eight adjacent COT cells in a given superlayer(each cell has 12 wires). A Mezzanine module receives those 96 wire inputs through the TDC module. The data(hits) on the 96 wires is compared to precision reference timing signals generated from the 132ns CDF_CLK. The Mezzanine identifies the data as either prompt or delayed bits for each wire. These bits are driven off the Mezzanine modules through the VME J3 backplane and onto the transition board as single-ended TTL levels every 22ns, i.e. 6 times the basic 132ns clock rate. In addition control bits such as Word zero, Beam zero, and a data strobe are sent.

The design of the XFT TDC Mezzanine module follows the guidelines specified in CDF note 4021, “Developing Mezzanine Cards”[4]. A block diagram of the XFT TDC Mezzanine module is shown in figure 2.

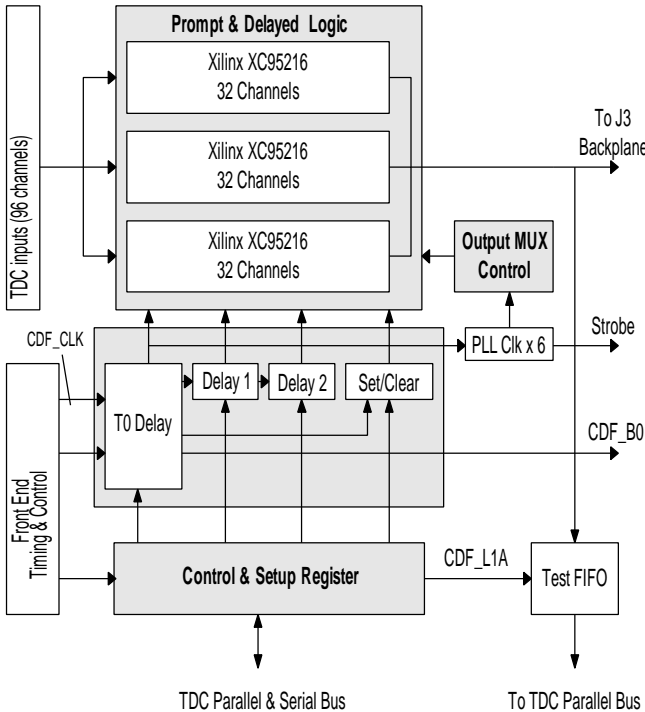


Figure 2: Block diagram of XFT TDC Mezzanine module.

A. Mezzanine Module Logic

The Mezzanine module operation can be broken down into three stages: 1) Create precision timing signals(gates) that can be used to compare against the cell data timing. 2) Classify the data as prompt or delayed with respect to the timing gates. 3) Multiplex and format the resulting data into 6 time bins so that it can be transported to the next stage of the system.

The precision timing signals are created by passing the 132ns CDF_CLK signal through a series of programmable silicon delays, specifically the Data Delay 3D7408S. The Data Delay 3D7408S is adjustable in 1ns increments. After the T0 delay the rising edge of the delayed CDF_CLK should be in synch with the earliest possible time at which valid cell data from the current crossing can arrive. The CDF beam crossing and CDF beam zero signal are also delayed to stay in synch with the delayed clock. The values for all the delays are kept in 8-bit registers in the ‘Control and Setup Register’ FPGA.

The delayed CDF_CLK is piped through Delay1 and Delay2 which provide clock edges that represent the end of the PROMPT and NOTSURE timing. The nominal widths of the PROMPT and NOTSURE windows are each 44ns. The “Control & Setup Register” FPGA uses these delayed clock edges to create 3 signals: “Prompt”, “Notsure” and “Late” that are in a logic high state during the various timing windows.

The delayed CDF_CLK is also piped to two other programmable delays, which are used to create the set and clear bits within the TDC circuit. The prompt and delayed set signals will come about 70ns and 140ns, after the delayed CDF_CLK edge.

The logic for an individual wire input, “Prompt & Delayed Logic”, is shown in figure 3. Precise timing of the XTC signals is critical for the proper performance of the XFT. To acquire this precise timing the “Prompt & Delayed Logic” is implemented in Xilinx XC95216 FPGAs which provide very tight timing stability. Our tests indicate that signal propagation times for the different channels within one chip vary by less than 1.5ns.

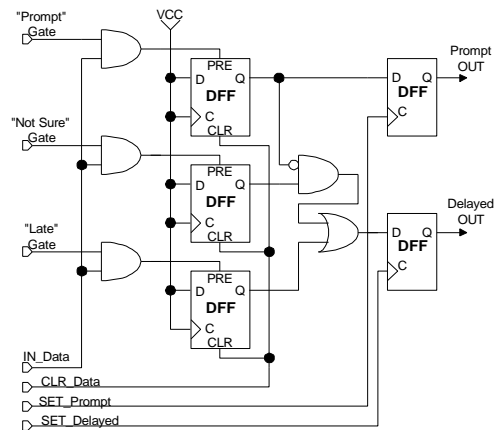


Figure 3: Mezzanine module Prompt & Delayed Logic.

Prompt And Delayed bits are multiplexed and sent off the Mezzanine module to the transition module in six 22ns cycles beginning with the next CDF_CLK after the current crossing. The 22ns clock is derived by running the CDF_CLK through a phase-locked loop device that is capable of producing a signal six times the input frequency. The Phase-locked loop device is a Motorola MC88915. The 22ns signal is sent to “Control & Setup Register” FPGA which generates 6 output enable signals. The output enable signals direct the correct 32 bits of prompt or delay data to the output pins. The 32 bits in the form of two 16-bit groups are transmitted off the Mezzanine module to a transition module.

On each cycle, we send to each Ansley cable the 16 data bits along with a strobe, Word zero marker, and Beam zero marker if appropriate. The prompt bits are sent on the first three cycles and the delayed bits on the last three. This protocol is shown in Table 1. The strobe runs continuously and is coincident with the data. The Word zero bit is present with the first 22ns slice of data and runs continuously. Beam zero is set for all six time slices of a “beam zero” event. The multiplexing scheme requires that COT wires 0-15 and 48-63 all be processed in one “Prompt & Delayed Logic” FPGA, wires 16-31 and 64-79 in the second, and 32-47 and 80-95 in the third.

Table 1
Mezzanine Module Output Format for 1st Ansley cable,
2nd Ansley cable handles wires 48-95.

Time Slice	Pairs 1,2,4,6, 24,25	Pair 3	Pair 5	Pair 7	Pairs 8-23
0	Gnd	Strobe	Beam 0	Word 0	Wires 0-15 Prompt
1	Gnd	Strobe	Beam 0	low	Wires 16-31 Prompt
2	Gnd	Strobe	Beam 0	low	Wires 32-47 Prompt
3	Gnd	Strobe	Beam 0	low	Wires 0-15 Delayed
4	Gnd	Strobe	Beam 0	low	Wires 16-31 Delayed
5	Gnd	Strobe	Beam 0	low	Wires 31-47 Delayed

IV. MEZZANINE TO FINDER TRANSMISSION

The Mezzanine to Finder link consists of the following pieces of hardware:

- COT TDC Transition modules, Quantity - 168
- Ansley cables, Quantity - 360
- Finder Transition Modules, Quantity - 60
- Custom VME J3 backplane, Quantity - 3

The COT TDC transition module receives single ended TTL level prompt and delayed data from the Mezzanine module. The transition module converts it to Low Voltage Differential Signaling (LVDS) technology and drives the data directly onto two Ansley cables.

The Ansley cable is a 200 foot flat cable with 25 differential signal channels. There are three wires per channel consisting of a balanced pair of adjacent signal wires and a ground wire which provides isolation from the next

channel. The cable has characteristic impedance of about 125 ohms. The rise-time is less than 7nsec (10% to 50% pulse height) and the cross-talk is less than 3%. The cable delay channel to channel varies less than +/- 1nsec.

Up to eight Ansley cables will be plugged into each Finder transition module. The transition module contains receivers, which convert the LVDS signals back to single-ended TTL signals, followed by buffers which reshape the single-ended signals. The custom J3 VME backplane directs the reshaped prompt and delayed data from the Finder Transition module into the Finder Module. The backplane will act to feed through all “core” Ansley cables and also to provide for the transmission of “neighbor” data from slot to slot if needed. Neighbor data must travel at most one slot across the backplane.




Figure 4: Picture of COT TDC Transition module, Ansley Cable, Finder Transition Module.

V. FINDER MODULE DESIGN

The Finder is designed to look for valid track segments in a given COT axial superlayer. Architecturally, the Finder Modules have been broken down into two types of modules. Each type of Finder Module will span 15 degrees of the COT. The SL1/3 Finder module will contain the logic for dealing with COT Axial Superlayers 1 and 3. The SL2/4 Finder, will contain the logic for dealing with COT Axial Superlayers 2 and 4. Figure 5 is a block diagram of the Finder SL2/4 module, the block diagram of the Finder SL1/3 module is similar.

The main logic of these modules will reside in the Finder circuits. Axial superlayers 1, 2, 3 and 4 will be instrumented with 48, 72, 96 and 120 Finder circuits respectively. Each Finder circuit is implemented in an Altera CPLD device. Each Finder CPLD receives its core input from a single Ansley cable and some additional information from “neighbor” cables. The Finder CPLD outputs track segment information that is passed to Linker Modules. A valid track segment is found whenever a predefined number of wires in a given 12 wire set(mask) have hit information on them. A valid segment is identified by 12 pixels for Axial superlayers 1 and 2 or 6 pixels and 6 bits of slope information for Axial superlayers 3 and 4.

A SL1/3 Finder Module will contain two SL1 Finder CPLDs and four SL3 CPLDs. There will be a total of 24 of these modules which will be spread out over 3 crates, with a total of 8 per crate. The custom J3 VME backplane requires that SL1/3 Finder modules be located in slots 4-11

A SL2/4 Finder Module will contain three SL2 Finder CPLDs and five SL4 CPLDs. There will be a total of 24 of these modules which will be spread out over the three Finder crates for a total of 8 per crate. The custom J3 VME backplane requires that SL2/4 Finder modules be located in slots 13-20.

A. FINDER MODULE Logic

The Finder module logic can be broken down into two sections: Control Logic and Data Flow. The data flow through the board starts with the Input and Alignment section followed by the Finder section and finally onto the Pixel Data Transmission section. With the implementation of reprogrammable devices there exist many diagnostic methods for testing the Finder module as a single unit or within the XFT system. A number of diagnostic designs have been made for the Alignment, Finder and Pixel section. The diagnostic designs in general allow the reprogrammable parts to be used as drivers or receivers with internal read/writeable RAM blocks that are capable of holding data to be driven or received.

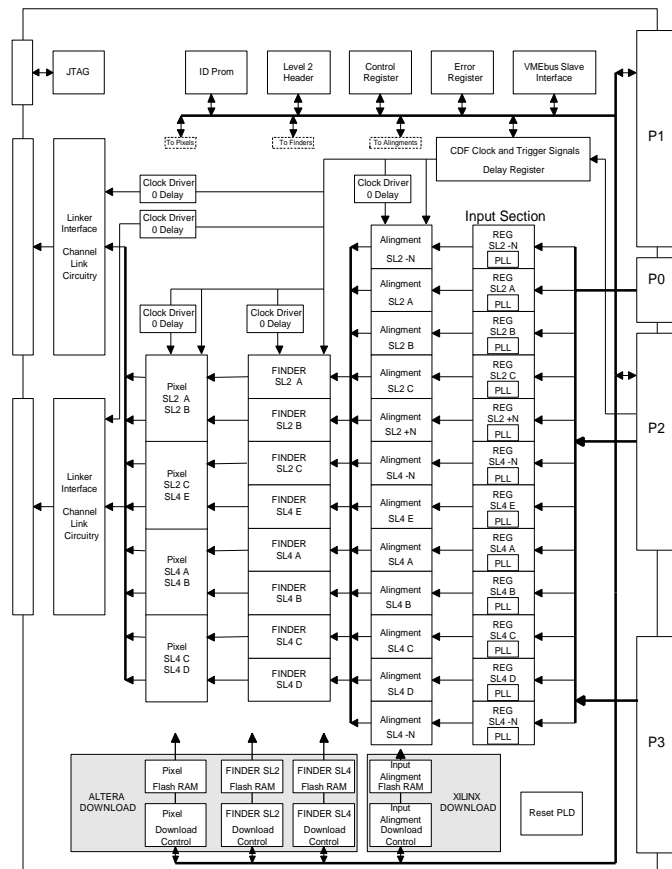


Figure 5: Finder SL24 Module Block Diagram.

• Control Logic

Control Logic includes: Clock circuitry, Flash RAMs, and various FPGAs that perform the tasks of communicating with the VME system for control and error monitoring. These blocks are all implemented in Xilinx 4000E series FPGAs. They are configured with the use of a serial EPROM on when power is applied or the reset button is engaged. The Finder module also implements a JTAG interface and Boundary Scan chain to provide a method for testing for infrastructure and interconnect defects.

It is possible to control the phasing of all clock and trigger signals on the Finder module. This is necessary to insure proper alignment of the raw data. The Finder board clocks are all derived from the 132ns CDF_CLK signal found on the VME backplane. The use of a programmable skew clock buffer specifically the Cypress CY7B991-7JC or Robo Clock allows for clock multiplication to derive a 33ns and 66ns clock.

The FINDER module will implement a modified version of a VMEbus slave interface. Only 32 bit aligned data transfers will be supported; these may be either single word transfers or block transfers. Only extended (32-bit) addressing modes will be supported. All modules will be assigned a unique geographical address through use of backplane pins on VME64extension backplane. FINDER modules will respond to the following address modifier codes: 09, and 0B. The VME_SLAVE design is implemented in a Xilinx XC4013E-3PQ240 FPGA.

The CONTROL_REGISTER provides a 32 bit register that is used to control functions in the Alignment, Finder and Pixel chips. The main functions are reset, loop and download.

The ERROR_REGISTER is used to register and count the Word_0 and Beam_0 errors generated in the Alignment and Finder chips.

The LEVEL 2(L2) HEADER WORD block provides a means to identify the Finder board and also the time at which a level 1 accept occurred with respect to the CDF Beam Zero signal. L2 accepts a hard coded board type, serial number and geographical address to identify the board. An 8 bit R/W register identifies the Finder boards pipeline depth. When a LEVEL 2 buffer is read from the Finder board the first 32 bit word that is read comes from the L2 FPGA and is used as a 'Header' for that LEVEL 2 data.

The ALTERA Download block consists of three Xilinx FPGAs which act as controllers for three Flash RAMs. The three sets are used to configure the Altera chips. The design provides a means to control the Flash RAM.

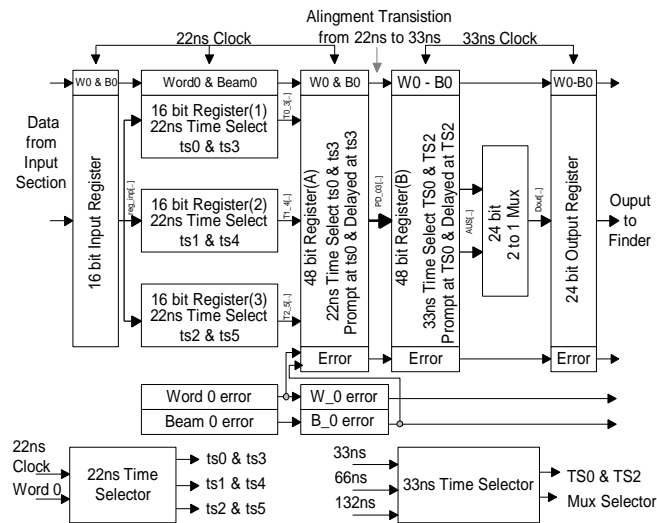
The XILINX Download block consists of a Xilinx FPGA which act as controller for a Flash RAM. The data in the Flash RAM is used to configure the 12 or 14 ALIGNMENT FPGAs.

The Flash RAMs on the Finder board are used to hold the Finder, Pixel and Alignment chips designs. The Flash RAMs are AMD AM29F040B-90PC or AM29F080B-90SC devices.

- *Data Flow*

Since the Ansley cables are presenting data every 22ns, a conservative approach has been taken, and an “input” stage will be used to receive the output of each Ansley cable. This input stage will make use of data registers that will be capable of 100 MHz synchronous operation, and phase lock loop(PLL) devices to regenerate the 22ns clock signal. The data registers used are the Cypress CY74FCT162823T 18-bit registers and the PLL are the Cypress CY7B991-7JC or Robo Clock devices. The input data from each Ansley cable will be registered by a signal formed by taking the edge “Strobe” signal coming from the same Ansley cable, running it through an individual Robo Clock, doubling its frequency, and allowing for phase adjustment. The registered data along with the regenerated 22ns clock is forwarded to the Alignment section.

The Alignment function is implemented in 12(Finder SL2/4 modules) or 14(Finder SL1/3 modules) Xilinx XC4005E FPGAs. These reprogrammable devices implement the design shown in the block diagram of Figure 6. Each FPGA is responsible for aligning the data from an individual Ansley cable to the 33ns onboard clock. The format of the data after it is aligned is shown in Table 2.



The Alignment FPGA design works by allowing the various 22ns registers to be used as storage areas while the data is being transferred to another block of registers operating off the 33ns clock system. The design works since the two clocks are a fraction of the 132ns CDF_CLK system. In general three time slices of the 16-bit data arriving at the 22ns rate are registered and held for 66ns. During that 66ns those 48-bits are transferred to another 48 bit register

Table 2.
Data Format as it leaves the Alignment FPGA.

time slice 0	Wire 0-23 Prompt, Word_0, Beam_0, Error
time slice 1	Wire 24-47 Prompt, Word_0, Beam_0, Error
time slice 2	Wire 0-23 Delay, Word_0, Beam_0, Error
time slice 3	Wire 24-47 Delay, Word_0, Beam_0, Error

The Finder section of the XFT system performs the job of identifying track segments in a given Axial superlayer of the COT. Finder circuits flag “hits” by setting pixels which indicate the position and/or slope of an identified track segment. Each Finder will report 12 pixels for each COT cell. A “hit” is identified to have occurred whenever at least 9, 10 or 11 out of 12 wires in a mask have been hit

```

graph LR
    Input[From Alignment Chips  
24 Core bits & 11 Neighbor bits @33ns] --> TD[Time Demultiplexer  
140 bits output @132ns]
    TD --> M41[4-1 Multiplexer  
48 bits output @33ns  
24 core bits output @33ns]
    M41 --> MB[Mask Block]
    subgraph MB [Mask Block]
        M1[4 input # of miss counter @33ns]
        M2[Sum of 3 "4 input # miss counters" @33ns]
        M3["OR" of Sum of 3 counters & Register @33ns]
    end
    MB --> Pixel[To Pixel Chip  
Pixel[11..0]]
    Pixel --> L1[Level 1 Wire Pipeline]
    L1 --> L2[Four Level 2 Buffers]
    L2 --> VME[VME Readout & Control Register]
    VME --> TD
    VME <--> DW[Dead Wire Reg.]
    VME <--> VMEBus[VME Data Bus]
    VME <--> AC[Address & Control]

```

The diagram illustrates the VME readout system architecture. It begins with data from alignment chips (24 core bits and 11 neighbor bits at 33ns) entering a Time Demultiplexer, which outputs 140 bits at 132ns. This data then passes through a 4-1 Multiplexer, which outputs 48 bits at 33ns and 24 core bits at 33ns. The output of the multiplexer enters a Mask Block, which contains three sub-components: a 4 input # of miss counter (33ns), a Sum of 3 "4 input # miss counters" (33ns), and an "OR" of Sum of 3 counters & Register (33ns). The output of the Mask Block is sent to a Pixel Chip, which outputs Pixel[11..0]. The Pixel Chip's output is then processed by a Level 1 Wire Pipeline, followed by Four Level 2 Buffers, and finally the VME Readout & Control Register. The VME Readout & Control Register has a feedback path to the Time Demultiplexer and is also connected to a Dead Wire Reg. and has bidirectional connections to the VME Data Bus and Address & Control lines.

The Finder CPLD operates as follows: The four time slices of data shown in Table 2 are received every 132ns along with pertinent neighbor cell information. That data is time demultiplexed into a 140 bit register that contains the prompt and delayed wire information. Those 140 bits are transferred to a multiplexer that selects a group of them(the

group size is superlayer dependent) every 33ns to forward to a Mask set. The Mask set consists of a large number of 12 bit masks(the number of masks is superlayer dependent). Within each mask the number of misses is counted. A miss is a wire without prompt or delayed information. If the number of misses is 3 or less, a pixel that relates to a segment in that cell is turned on. There are separate MASK designs files that look for 1, 2 or 3 misses for each of the superlayers. The pixel information is registered every 33ns and transferred to the output pins.

The Finder CPLD also contains a Level 1 pipeline and four Level 2 buffers that store prompt and delayed wire information for VME readout.

4) Pixel Data Transmission

Each Finder module must present two copies of its pixel (or, pixel and slope) information to the Linker Modules. This task is performed using a CPLD called the Pixel Chip and Channel Link drivers. The Pixel chip duplicates and forwards two Finder chips worth of information to two separate groups of three(Finder SL1/3) or four(Finder SL2/4) Channel Link drivers which in turn forward the information to different Linker modules. The Channel Link devices(National Semiconductor's DS90CR281/282 28-bit pair) consist of a driver and receiver pair. The driver is located on the Finder board and the receiver is located on the Linker board. Channel Link devices are grouped together to reduce the number of connectors and cables between the modules.

The Pixel Chip designs are implemented in the Altera Flex 10K EPF10K20RC240-3 CPLD for superlayers 2, 3 and 4 and the Altera EPF10K50RC240-3 CPLD for superlayer 1. The Pixel chip performs three functions: 1) Combine the Pixel data from two Finder chips into a single 28 bit word. 2) Duplicate that information and drive it through separate individual output pins to individual Channel Link drivers. 3) Provide a Level 1 pipeline and four Level 2 buffers for Pixel data for VME readout.

VI. FINDER TO LINKER TRANSMISSION

The link between the 48 Finder modules and 24 Linker modules consists of 96 round cables and numerous National Semiconductor's DS90CR281/282 28-Bit Channel Link devices. Each Linker module receives pixel data from four separate Finder modules on four separate round cables. The Finder module drives pixel data to two different Linker modules on two separate round cables.

Each Finder module implements two groups of 3(SL1/3 modules) or 4(SL2/4 modules) National Semiconductor's DS90CR281 28-Bit Channel Link drivers grouped together. The DS90CR281 driver converts 28 bits of data into four LVDS data streams every clock cycle(33ns). This same clock is phase-locked and transmitted in parallel with the data streams over a fifth LVDS link. A group of eleven DS90CR282 Channel Link receivers on the Linker module convert the LVDS data streams back into TTL data in synch with eleven individual Channel Link output clock signals. The output clock signals are derived from the transmitted

LVDS clock signals thru a phase lock loop. The data on the LVDS pairs will be transmitted at 7 times the clock frequency. The LVDS clock signal is transmitted at the PLL frequency.

The cable run from Finder to Linker Modules will be six to ten feet in length. The cables are made out of 3M 3600 series cable and 3M Mini D Ribbon connectors.

A SL1/3 Finder module utilizes an 18 pair cable with 36 pin connectors and the SL2/4 Finder module utilizes a 25 pair cable with 50 pin connectors.




Figure 8: Linker module, Channel Link Cables and Finder Module.

VII. LINKER MODULE DESIGN

The Linker module has the responsibility of linking segments between the axial superlayers of the COT to form a track. The segments which are defined by pixels are found by the Finder modules. The Linker is separated into 288 phi slices of 1.25 degrees. Each slice has a dedicated chip with the task of finding tracks within the slice. The tracks found by the Linker chips are passed to the XTRP for extrapolation to other portions of the detector and used in the Level 1 trigger decision. A block diagram of the Linker module is shown in figure 9.

The Linker system will consist of 24 identical 9U VME modules each with 12 Linker chips. The modules reside in three Linker crates (8 cards in each crate) that are located above the Finder crates in the three XFT racks in the first floor counting room.

A. LINKER MODULE Logic

The Linker module as with the Finder module can be divided into two sections – Control logic and Data flow. The data flow through the Linker board starts with the receiving of pixel data by the Channel Link receivers. That received data is aligned to the board clock in the Input Formatter section which forwards the data to the Linker chips for track identification and finally onto the Output Formatter. A number of different diagnostic designs have been made for the Input Formatter, Linker and Output Formatter section. These diagnostic designs allow the reprogrammable parts to be used as drivers or receivers with internal read/writeable RAM blocks that are capable of holding data to be driven or received.

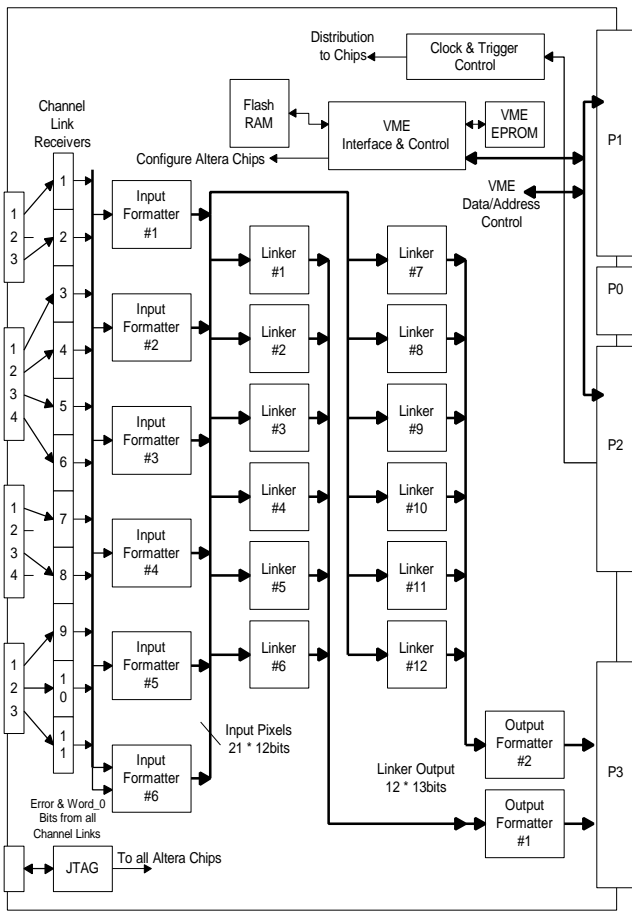


Figure 9: Linker Module Block Diagram.

• Control Logic

Control Logic includes: Clock circuitry, Flash RAMs, and a VME_SLAVE that performs the tasks of communicating with the VME system for control and error monitoring. The Linker module also uses a JTAG interface for testing and programming.

The Linker board implements a modified version of a VMEbus slave interface in a Altera Flex 10K EPF10K30RC240-3 CPLD. Only 32 bit aligned data transfers will be supported; these may either be single word transfers or block transfers. Only extended (32-bit) addressing modes will be supported. The interface controls the loading of programs into the on board Flash RAM along with the download of the Flash RAM data to the Formatter and Linker CPLDs.

The on-board 33nsec, 66nsec, and 132nsec clocks are all derived from the CDF_CLK. The on-board 132nsec clock is phase shifted from the backplane CDF CLK using the DS1020 silicon delay line from Dallas Semiconductor. This chip allows delays of up to 256 nsec, in 1 nsec increments. The delayed clock is then manipulated in a scheme similar to the Finder modules using the Cypress CY7B991-7JC or Robo Clock. The resulting clock signals are redistributed throughout the module with National Semiconductor's Low skew clock buffer CGS2534V.

• Data Flow

1) Input Formatter

There are 6 Input Formatter chips per board which are implemented in Altera Flex 10K EPF10K40RC240-3 CPLDs. The Input Formatters latch the data from the Channel Link receivers using the Channel Link output clocks, then synchronize the data to the onboard 33nsec clock. The output is distributed to the Linker chips. The Input chips also contain the pixel data level 1 pipeline and four level 2 buffers for VME readout. The 6th Input Formatter chip is also used for error checking: comparing word_zero, beam_zero and start bits from all Channel Link receivers.

2) Linker chips

The Linker chips are implemented in Altera Flex 10K EPF10K50RC240-3 CPLDs. Each Linker accepts information from 6 Input Formatter chips, and searches for the best track. There are 12 Linker chips per board, each covering 1.25 degrees for a total coverage of 15 degrees per module Figure 10 shows a block diagram of the Linker chip.

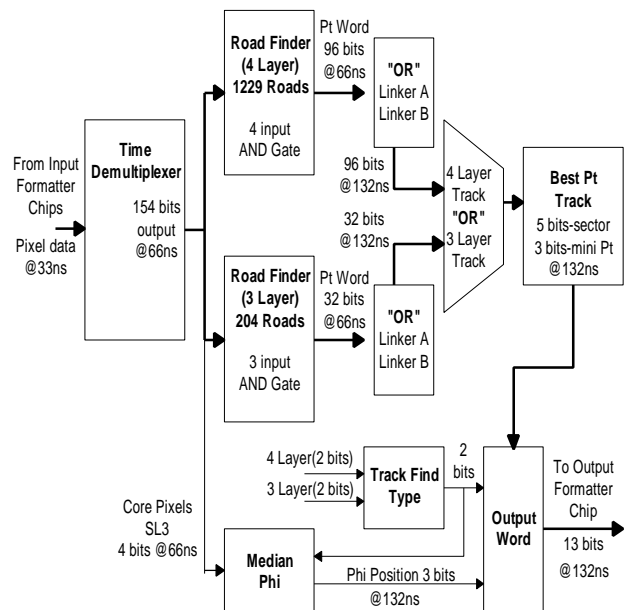


Figure 10: Linker CPLD Block Diagram.

The Linker chip operates as follows: Each finder sends the information from the 4 COT cells it processes to the Linker in 4 12-bit words, 1 word every 33nsec. These words then have to be time-demultiplexed so the linker can look at all cells in parallel, for all 4 superlayers. The inputs to the 12 Linker chips are slightly different, and so this routine is different for each Linker chip. There is a two-fold symmetry in the Linker, which allows the number of stored roads to be cut in half. Effectively, each Linker is then arranged as two sub-linkers, which we call A and B.

The pixels for Linker A and B are passed on to the road finding logic. There are 1229 4-layer Roads and 204 3-layer Roads implemented in the chip logic. Each road is stored as

the ``AND" of 4 (or 3) pixels, one from each layer. Four-layer and three-layer roads are found in parallel. The four-layer road finding outputs a 96 bit Pt word every 66ns. Each bit in this word corresponds to a valid track found with a given Pt bin. The three-layer road finding outputs a 32 bit Pt word every 66ns. The number of output bits is smaller because the Pt resolution is diminished by using only 3 layers rather than 4.

Either the 96 (4 layer) or 32 (3 layer) bit Pt word is presented to the ``Best Pt Track" which determines what Pt to report. The 3 layer track will be presented if there isn't a 4 layer track. The BEST Pt logic is designed to report the median Pt bin of this cluster of Pt bins. In addition, the BEST Pt logic must handle the case where more than one track passes through a given Linker chip. The bias at present when this occurs is to select the cluster of Pt bins which is higher in Pt and then report the median of this cluster. The output of "Best Pt Track" is a 8 bit word with Pt and charge information: 3 bits of mini-Pt and 5 bits of section Pt .

The "Median Phi" block looks at the 8 pixels reported from the FINDER for superlayer 3 and determines a simple median of the pixels, outputting a 3 bit word. Remember that tracks are found in sub-linkers A and B. If no track is found in a given sub-linker, the pixels for superlayer 3 for that sub-linker are masked off.

The Linker outputs information on the best track found every 132nsec. The output is 8 bits of Pt , 3 bits of phi, and 3 bits of code. The eight Pt bits are 3 bits of mini-Pt and 5 bits of section Pt . Only 7 bits are required to uniquely specify a Pt, but for space reasons this conversion is done in the Output Formatter chip.

3)Output Formatter

There are 2 of these chips per board implemented in Altera Flex 10K EPF10K30RC240-3 CPLDs. Each chip inputs 14 bits of data from 6 Linker chips every 132nsec. The Output Formatter forms the 7 bit Pt word using the 3 bit Mini Pt and 5 bit Section Pt from the 6 Linkers. In addition the Output Formatter implements the level 1 pipeline containing 78 bits(12 bits per Linker chip) of Linker information(Pt, Phi, Track isolation, and whether track was 3 layer or 4 layer).Duplicate tracks near the boundaries of the region covered by each Linker chip are removed and the tracking information is multiplexed into 4 data words that are sent out sequentially (every 33ns) to the XTRP system.

VIII. LINKER TO XTRP TRANSMISSION

The found track information provided by the Linker module is transmitted one slot across the VME backplane to a Linker Transition module which converts the single ended signals to LVDS signals using National Semiconductor DS90CR31 devices. The LVDS drivers transmit the data across 50 feet of shielded twisted pair cable to a transition module in the XTRP crate. The 100 conductor twisted pair cable is produced by AMP inc. and uses the Amplimite .050 Subminiature D connectors. The transition modules use AMP 100 pin right angle Amplimite connectors.

IX. CONCLUSIONS

The XFT system is in the production stage with all modules either at the printed circuit board manufacturer or assembly contractor. Extensive testing has been performed to insure that all modules will perform the tasks detailed in this paper. Prototypes of the Modules and transmission medians have been tested alone and in a reduced system capacity. The installation date of the XFT system is projected to occur in the fall of 1999.

X. REFERENCES

- [1] Ken Bloom, et al., XTC: The COT TDC Mezzanine Card, CDF/DOC/Tracking/Public/xxxx, 1998.
- [2] T. Shaw, et al., Specification for XFT Finder Modules SL1/3 Module and SL2/4 Module, CDF/DOC/Tracking/Public/xxxx, 1999.
- [3] R. Hughes, et al., Run II XFT Linker Design Specification, CDF/DOC/Tracking/Public/xxxx, 1999.
- [4] M. Campbell, et al., Developing Mezzanine Cards, CDF/DOC/Tracking/Public/4021, 1998.
- [5] P. Koehn, et al., Online Track Processor for the CDF Upgrade, CDF/DOC/Tracking/Public/xxxx, 1999.